Systems-on-a-Chip (SoCs)

An introduction to SoCs using the Cell® processor as a case-study

Presented by Michael Fuller

What is a system-on-a-chip?

"An SoC is an ASIC that integrates, on a single silicon die, processors, memories, logic, analog, and I/O functions previously implemented as multiple discrete chips."

T. R. Bednar et al, "Issues and strategies for the physical design of system-on-a-chip ASICs." *IBM J. RES. & DEV.*, vol. 46, no. 6, Nov. 2002.

RF Micro Devices Bluetooth SoC

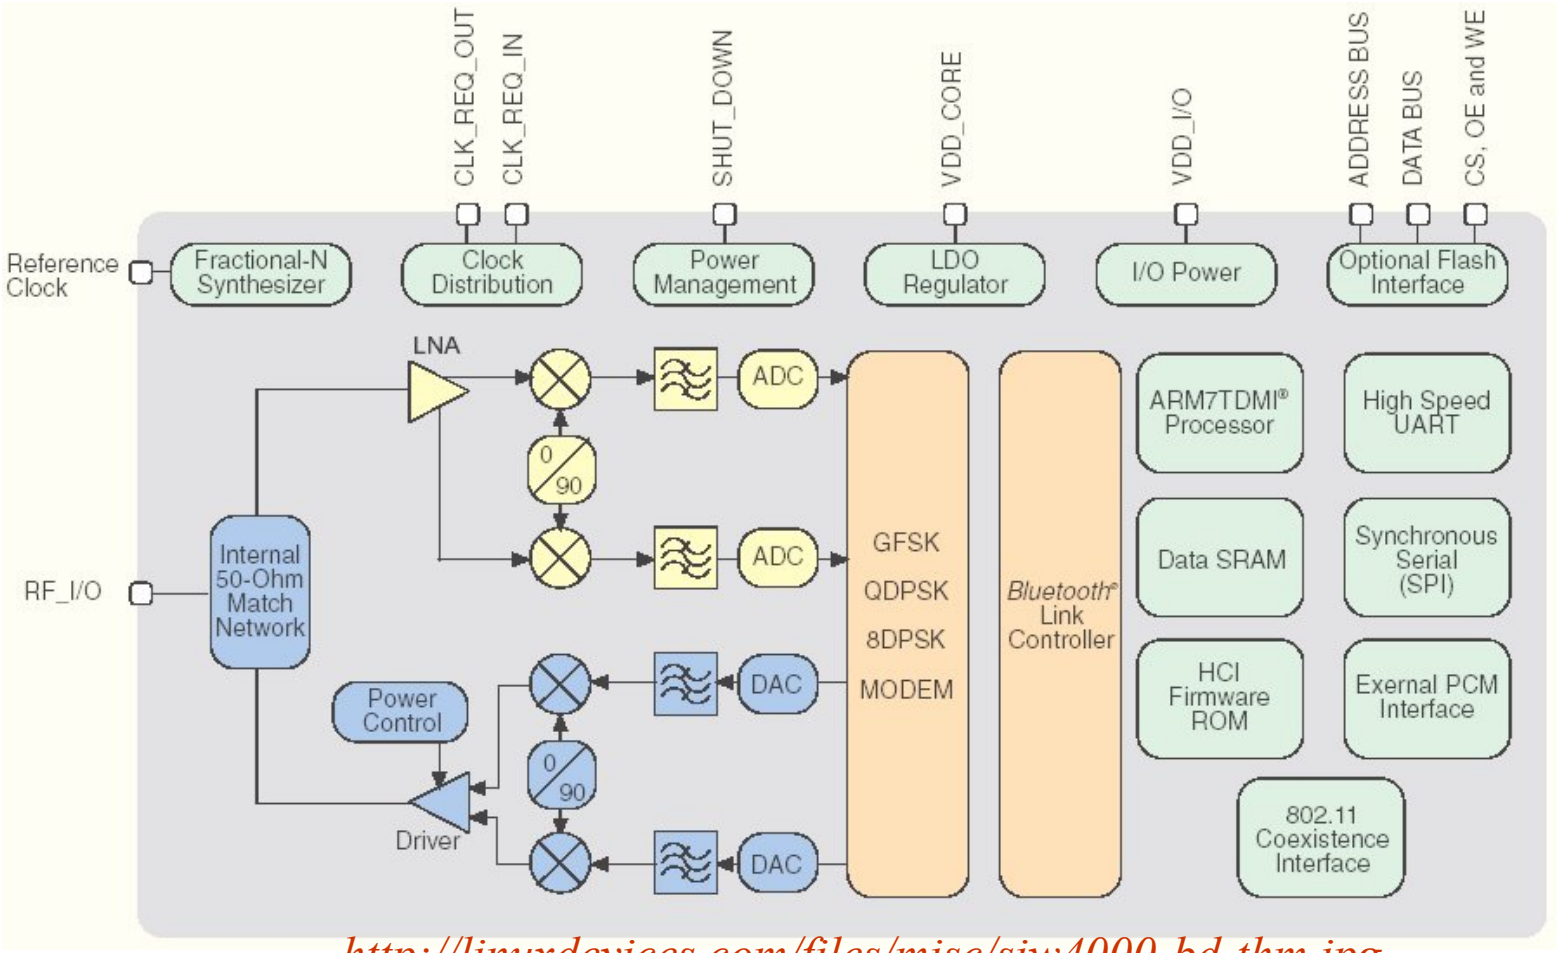

http://linuxdevices.com/files/misc/siw4000-bd-thm.jpg

Requires only 8 external components (6 caps, 1 inductor, 1 BPF)

Cypress PSoC™ Mixed-Signal Array

- Single chip containing programmable digital and analog function blocks
- Programmable interconnect
- Low Cost: $0.05 - $10
- 24 MHz, 4 MIPS, 8-bit µCPU
- Flash program memory
- SRAM data memory
- Configurable I/O

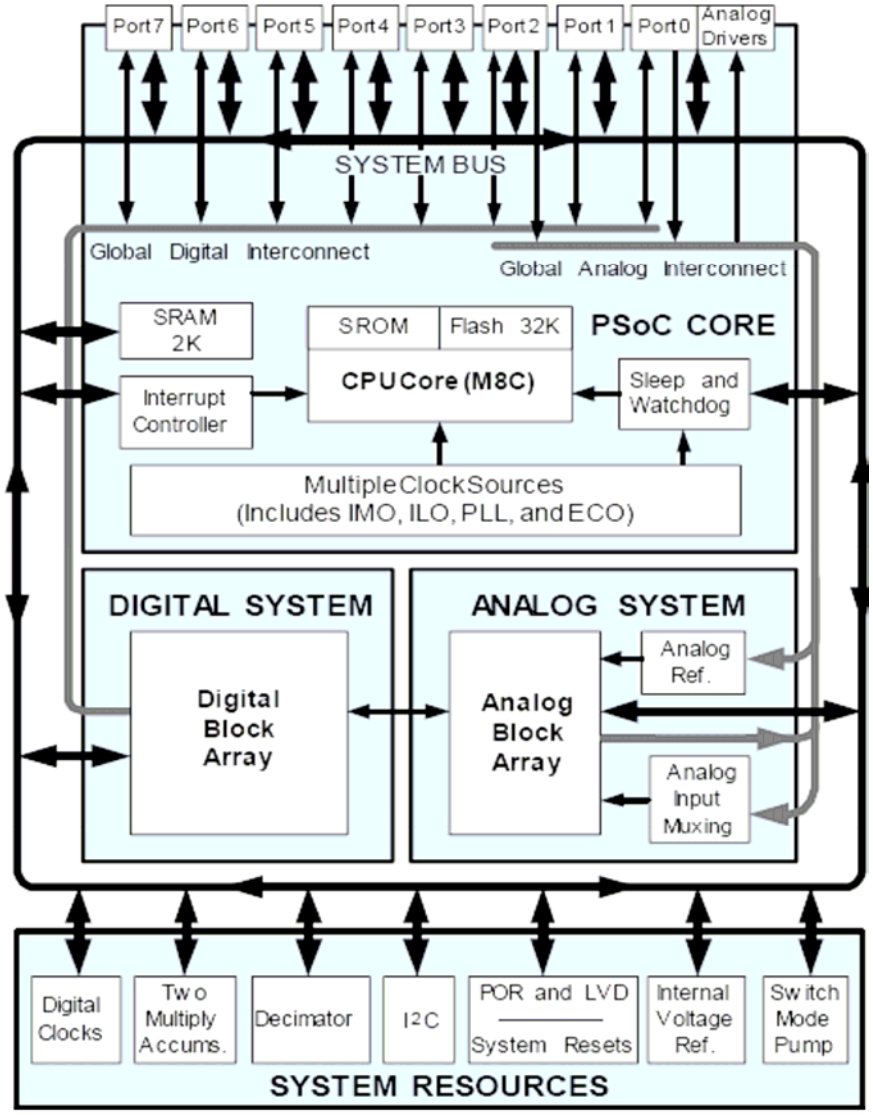

Cy8C29x66 datasheet, Cypress Microsystems, Inc., Nov. 12, 2004.

Cypress PSoC™ Mixed-Signal Array

- Digital System:
  - PWMs (8-32 bit)
  - Counters (8-32 bit)
  - Timers (8-32 bit)
  - UART w/ selectable parity
  - SPI master and slave (4 each)
  - Pseudo Random Seq. Generator
  - More...

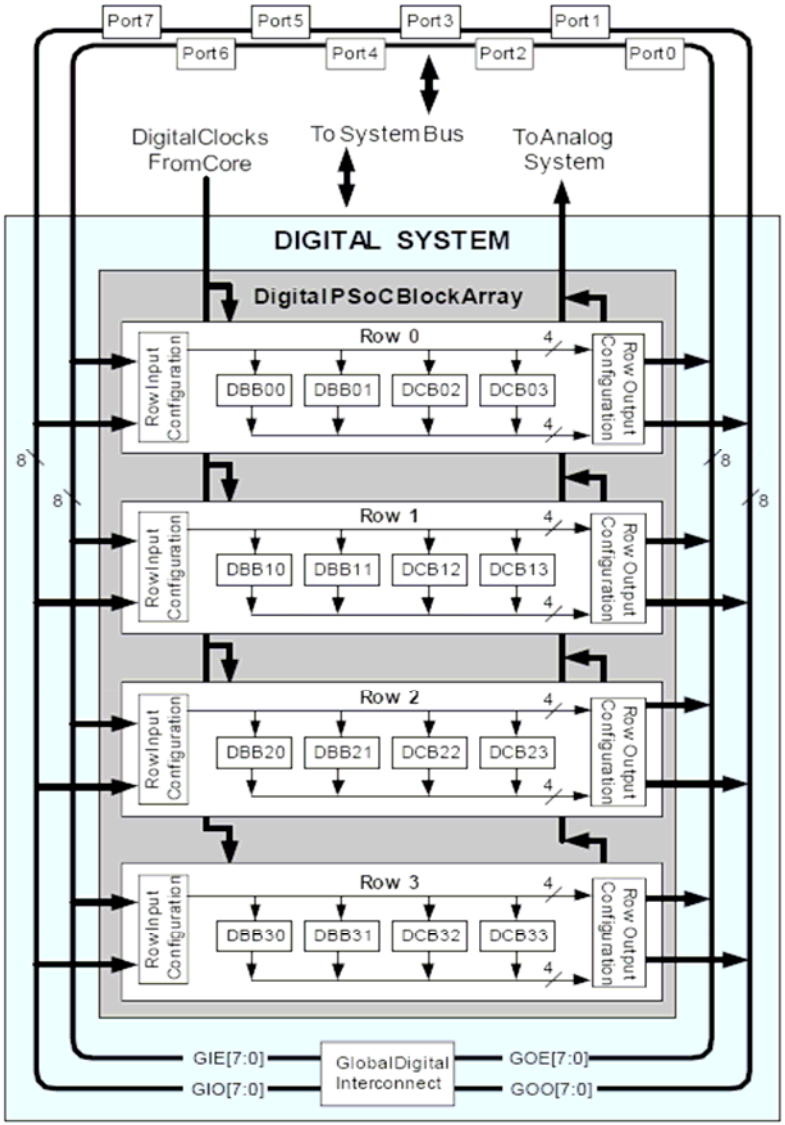

Cy8C29x66 datasheet, Cypress Microsystems, Inc., Nov. 12, 2004.

Cypress PSoC™ Mixed-Signal Array

- Analog System:
  - ADCs (6- to 14-bit resolution)
  - Filters (2,4,6,8 pole BP, LP, notch)
  - Amplifiers (selectable gain to 48x)
  - Instr. Amplifiers (sel. gain to 93x)
  - Comparators (16 sel. Thresholds)
  - DACs (6- to 9-bit resolution)
  - High current output drivers (40 mA)
  - 1.3 V reference
  - DTMF dialer
  - Modulators
  - Correlators
  - Peak Detectors
  - Many other topologies possible...

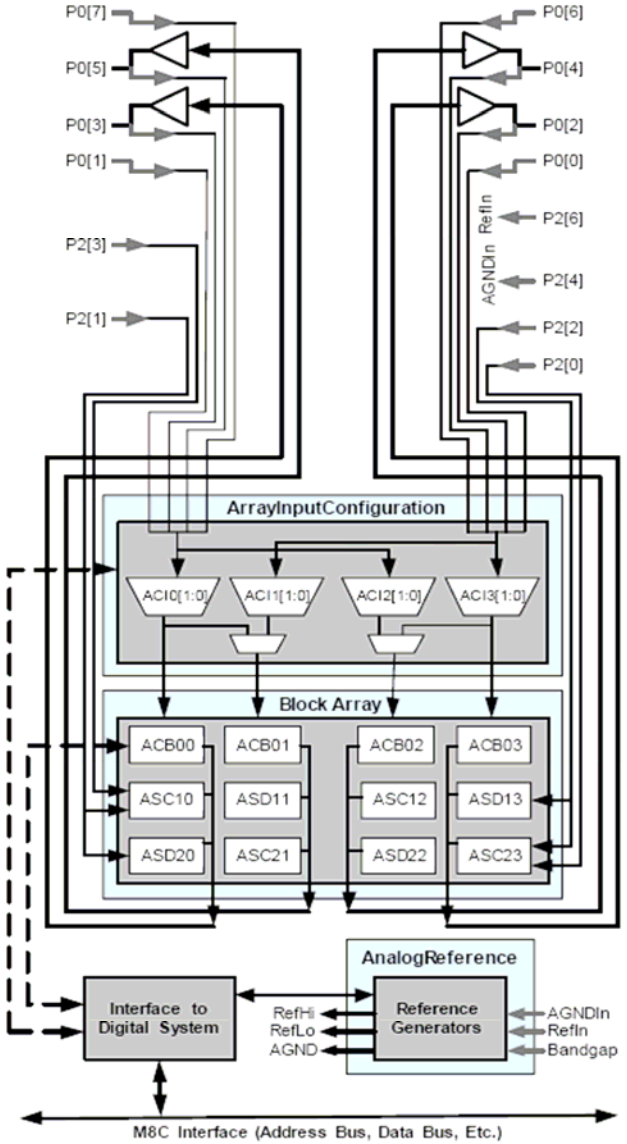

Cy8C29x66 datasheet, Cypress Microsystems, Inc., Nov. 12, 2004.

Issues for SoC Designs

- Semiconductor technology
  - Transistor type (CMOS, SOI, etc.)
  - Min. feature size (0.5 micron or 65 nm?)
  - Gate-oxide thickness
  - Leakage current
  - Vt (high, low, multiple)
  - Analog transistors (long-channels, thicker oxide)?
  - Analog process elements (polycaps, precision resistors)
  - Single or multiple well
  - Metal stack (# layers, thin/thick, ↑layers ➔ ↑cost)
  - Memory process features (e.g., trench caps)

Issues for SoC Designs

- Packaging
  - Flip-chip?
  - Wirebond?
  - Number of I/O pins?
  - Number and distribution of power pins?
  - Thermal management

Issues for SoC Designs

- Power
  - Digital Power
  - Analog Power
    - Must be isolated from noisy digital circuitry and supplies
  - Multiple supplies (multiple VDDs, bulk biases, etc.)
  - Routing issues
  - Power management!
    - Sleep mode, deep-sleep mode, standby mode, idle mode, low-power mode, high-performance mode, etc., etc.
    - Throttling, voltage scaling, frequency scaling

Issues for SoC Designs

- Memory
  - DRAM
  - SRAM
  - Flash
  - ROM, PROM, EEPROM
  - L1/L2 cache
  - data memory, instruction memory
  - Memory interface (bus architecture, control, DMA)
  - Analog memory

Issues for SoC Designs

- Clocking
  - On-chip or off-chip?
  - Multiple clocks?
  - Asynchronous function blocks?
  - Routing and distribution
  - Skew management
  - PLLs
  - Frequency dividers/synthesizers

Issues for SoC Designs

- Testing and Verification
  - Pre-Tapeout
    - Simulate huge design (100s of millions of transistors)
    - Simulate both analog and digital components (some interaction)
    - Noise sims, power sims, thermal sims
  - Post-tapeout
    - Test/Debug modes
    - Test/Debug hardware (extra buffers, config. registers)
    - JTAG
    - Calibration
      - Laser trimming
      - Configuration registers
      - Redundant memory blocks

The CELL processor

Joint effort by IBM Corp., Sony, and Toshiba

- CELL = Cell Broadband Engine Architecture
- Intended for concurrent real-time multimedia and conventional computing applications
- $400 million budget
- 400 engineers
- 90-nm SOI process
- 10 simultaneous threads on 9 processors
- 256 GigaFLOPS @ 4 GHz
- 235 million transistors
- 221 mm^2 chip area
- 60-80 watts @ 4 GHz

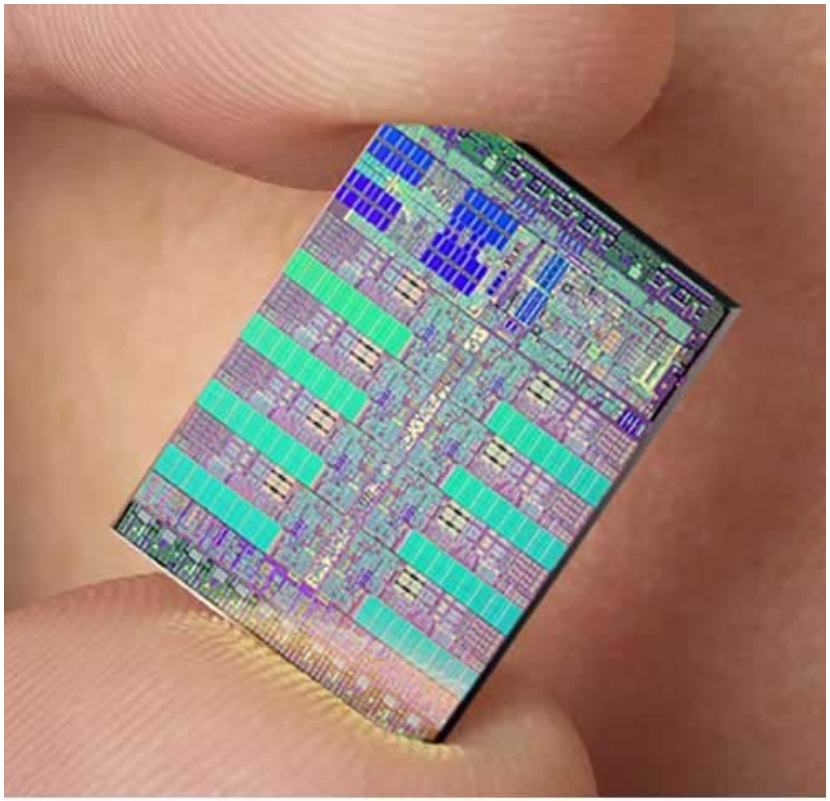

Implemented on a 90-nm silicon-on-insulator process with eight levels of copper interconnect, the CELL processor can deliver a single-precision compute throughput of over 256 GFLOPS when running at 4 GHz.

http://www.elecdesign.com/Files/29/9748/Figure_02.jpg

First intended application:

Sony PlayStation 3

Release Date:

- early November 2006

http://en.wikipedia.org/wiki/Image:ME0000570927_2.jpg

http://techon.nikkeibp.co.jp/english/NEWS_EN/20050425/104149/20050424toshiba_48.jpg

Toshiba demonstrates the Cell processor simultaneously decoding 48 SDTV format MPEG-2 streams in 2005.

Main Components of the CELL

- Power Processing Element (PPE)
- Synergistic Processing Elements (SPEs)
- Element Interconnect Bus (EIB)
- Memory Interface Controller (MIC)
- Bus Interface Controller (BIC)

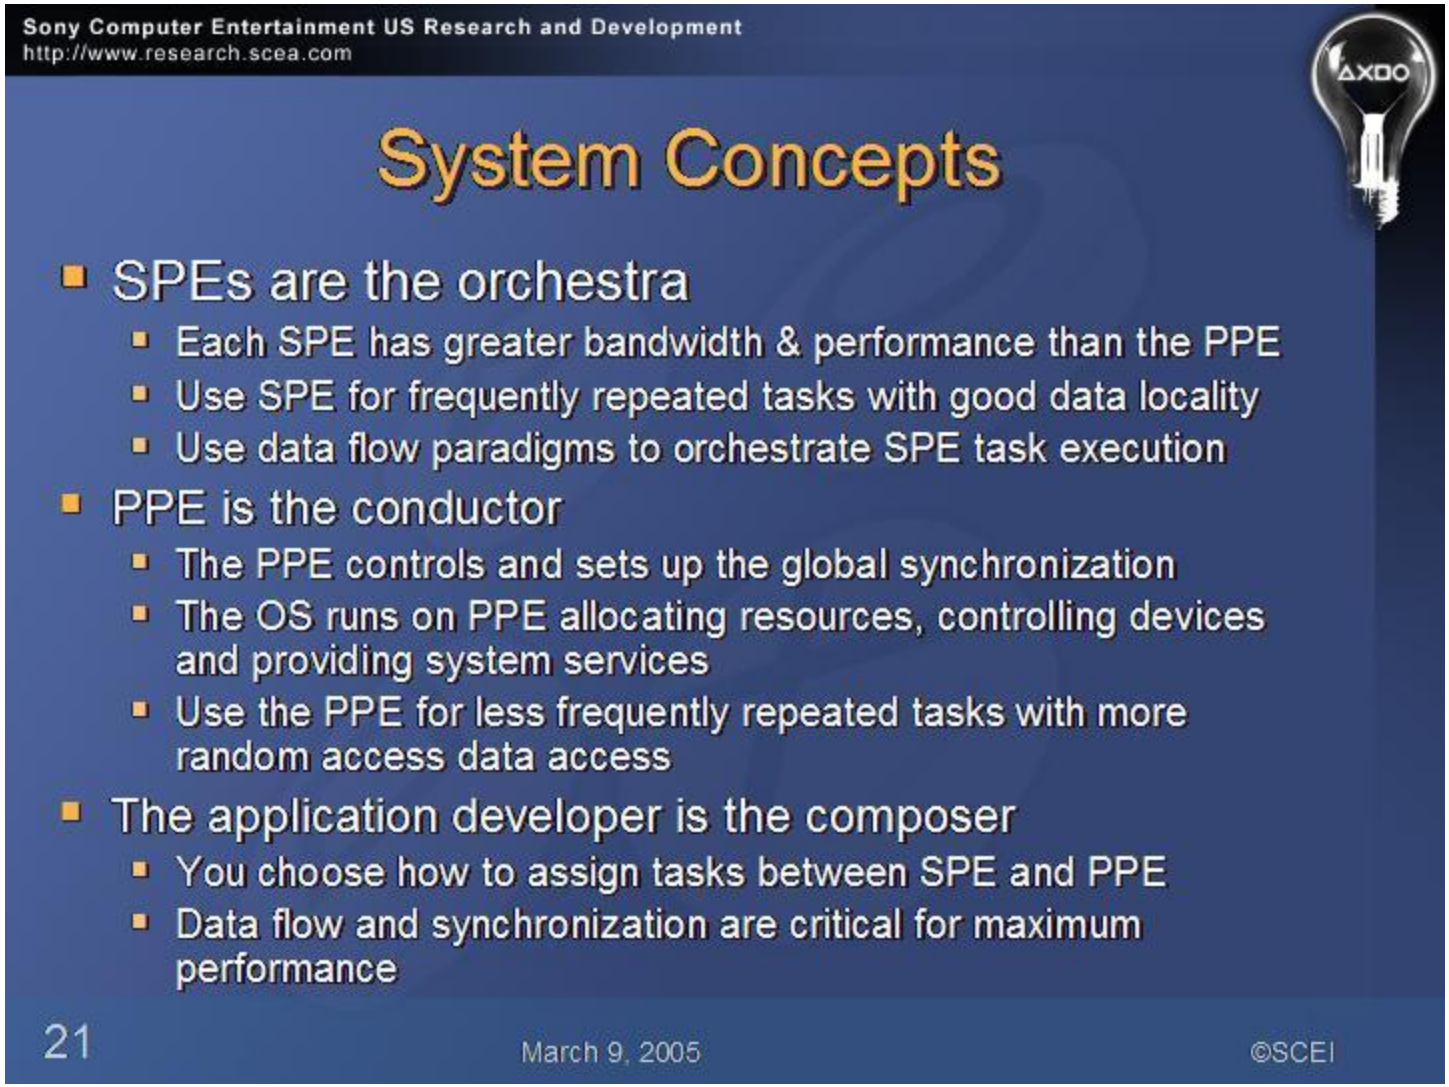

http://research.scea.com/research/html/CellGDC05/21.jpg

Cell Processor Architecture

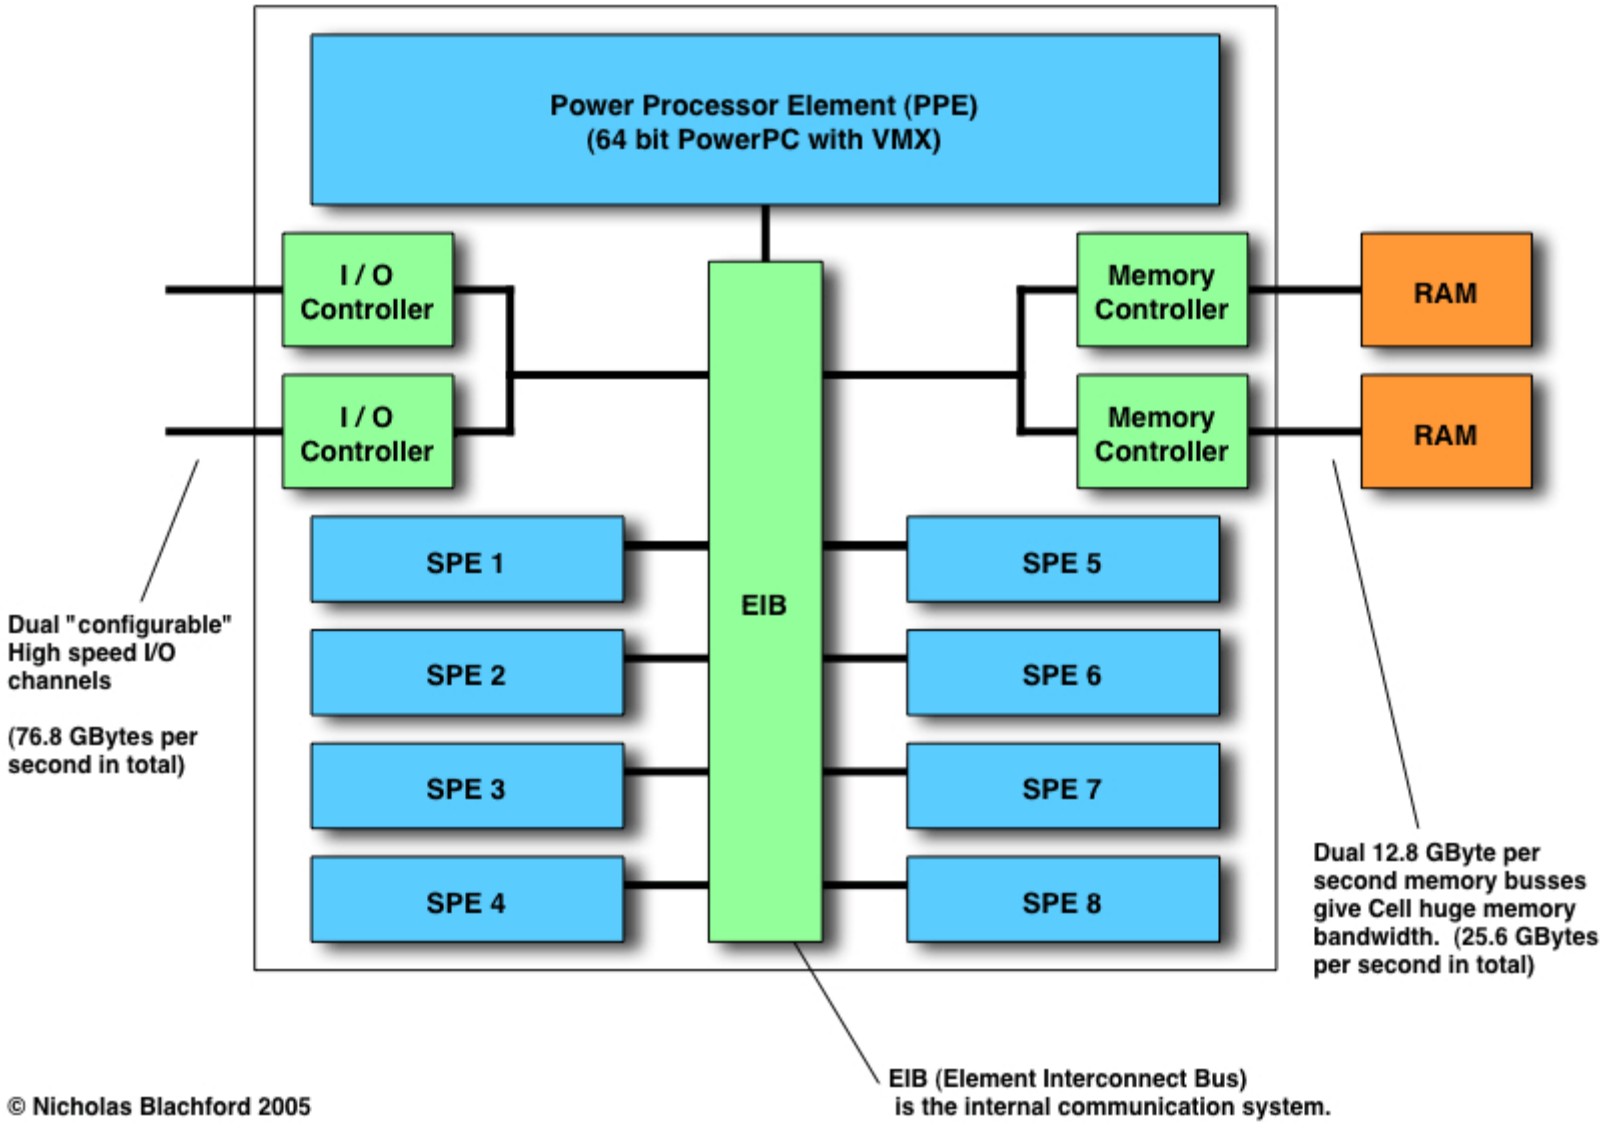

http://www.blachford.info/computer/Cell/Cell_Arch.gif

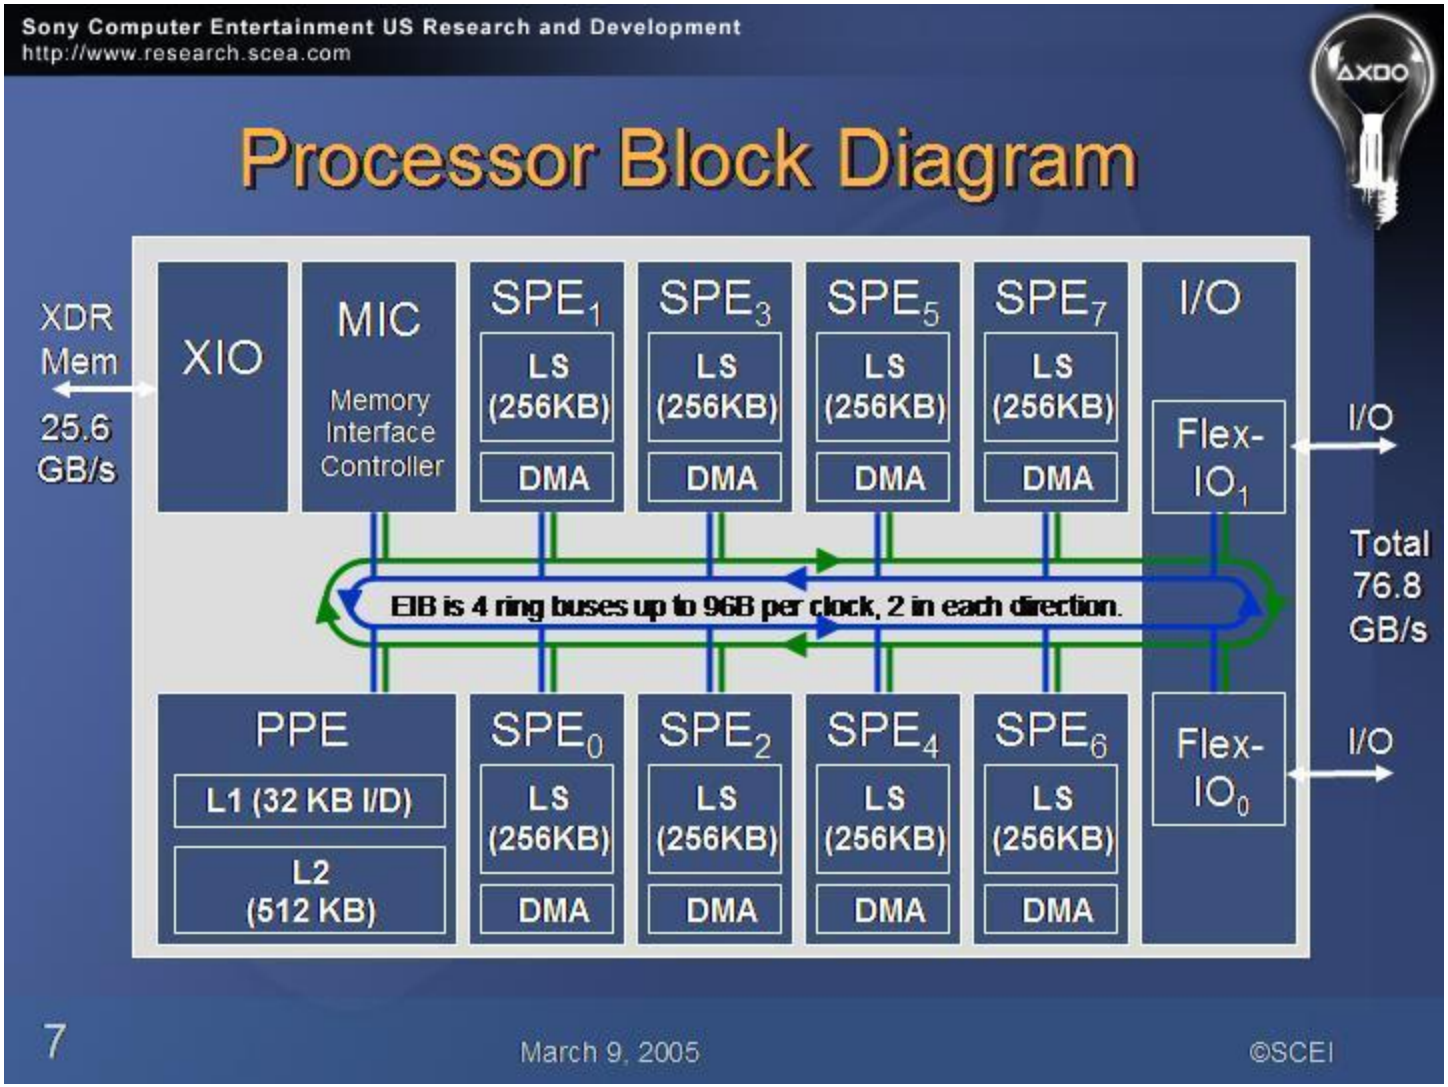

http://research.scea.com/research/html/CellGDC05/08.html

Cell Chip Layout

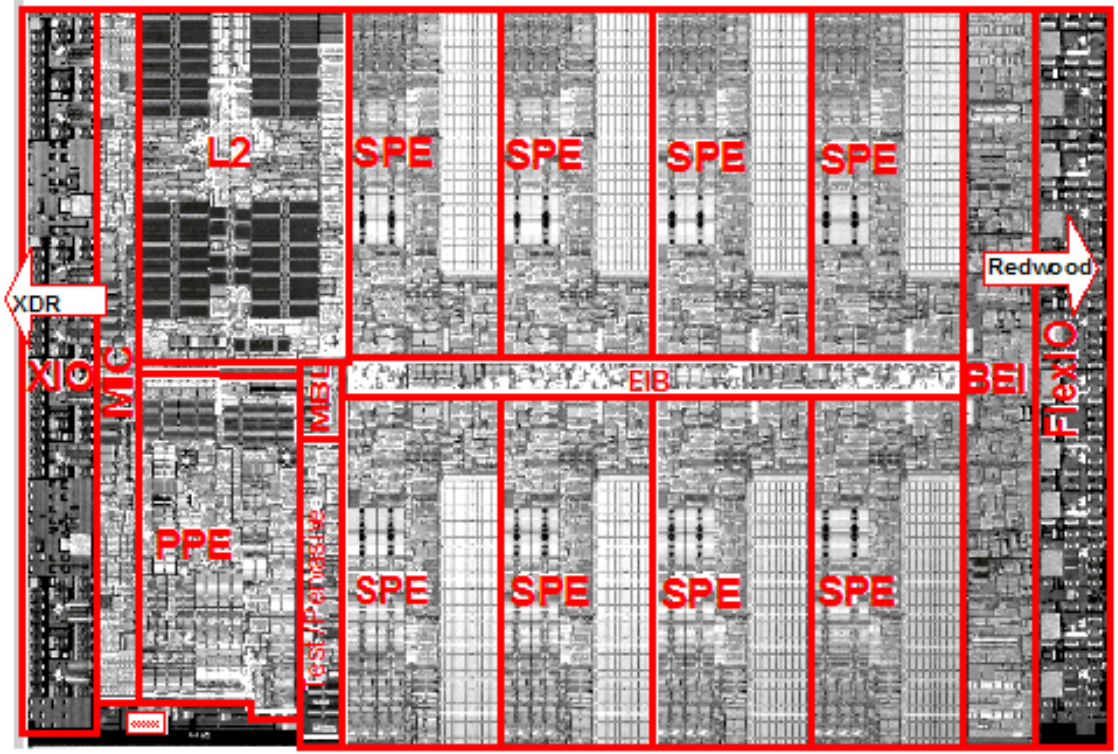

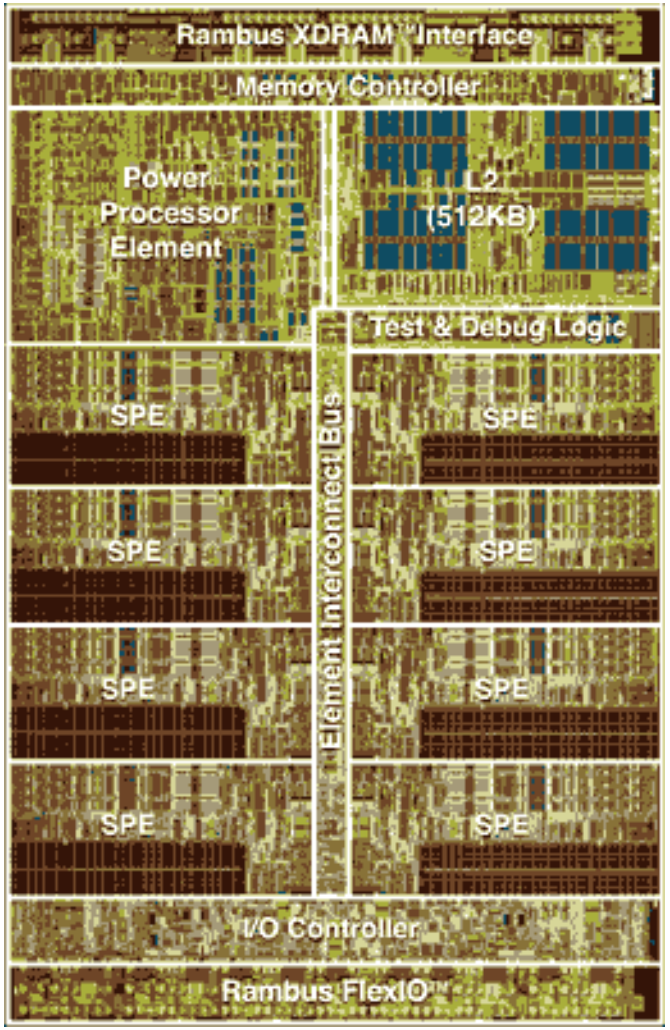

PPE

Power Processing Element

- Dual-threaded, 64-bit core
- 32-kB instruction/data L1 caches
- 512-kB L2 cache
- Power Architecture family compliance:
  - Integer unit
  - Floating point unit
  - VMX unit
  - MMU unit
- Handles DMA requests for SPEs
  - via memory-mapped IO controllers

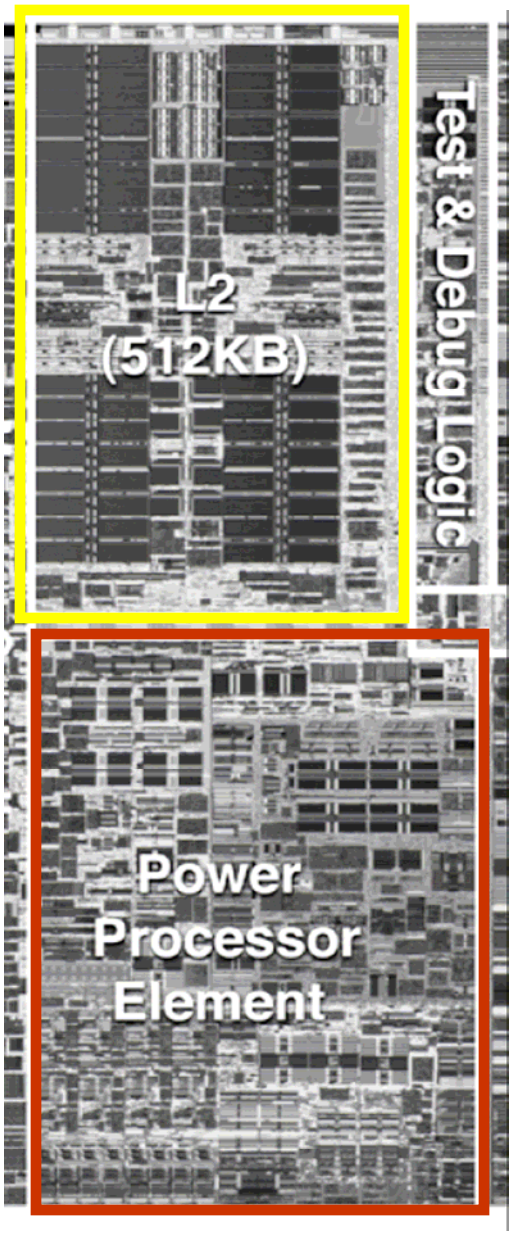

SPE

Synergistic Processing Element

- Data and instructions stored in private 256kB local storage (LS)
- SPEs share system memory (L2 cache) with PPE
- Asynchronous data transfers between LS and main storage (L2 cache)
  - Communication overlaps with computation ➔ real-time operation

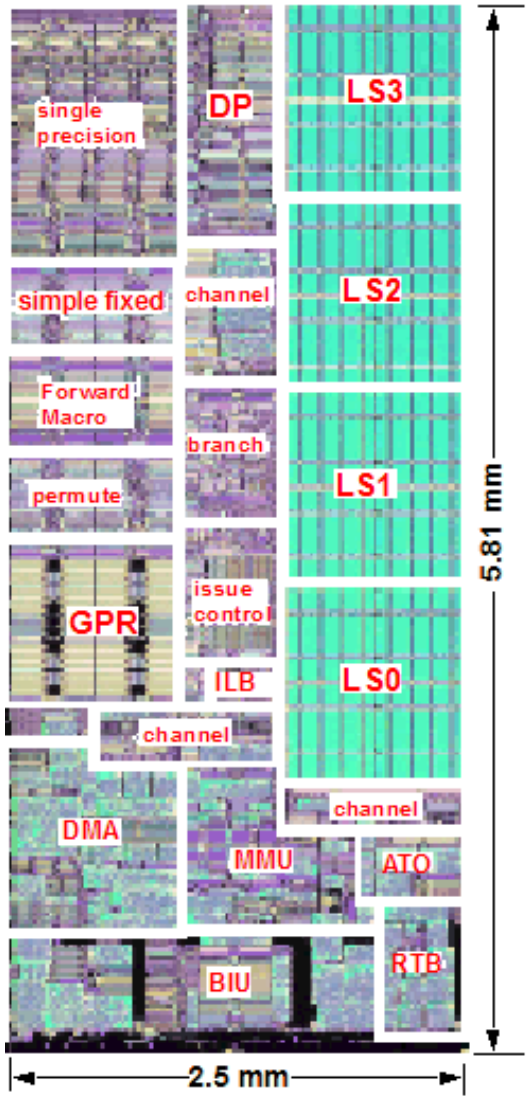

http://www.realworldtech.com/includes/images/articles/cell-3.gif

EIB

Element Interconnect Bus

- Four 16-byte-wide data rings
  - Each support 3 simultaneous data transfers
- 96 bytes per processor cycle

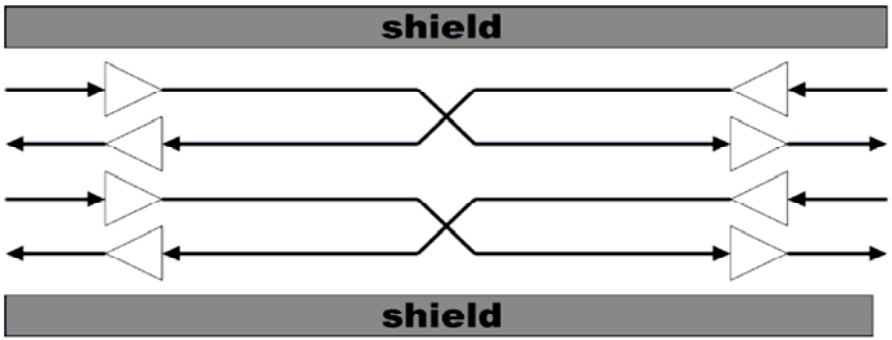

Physical interleaving of bus to minimize noise coupling

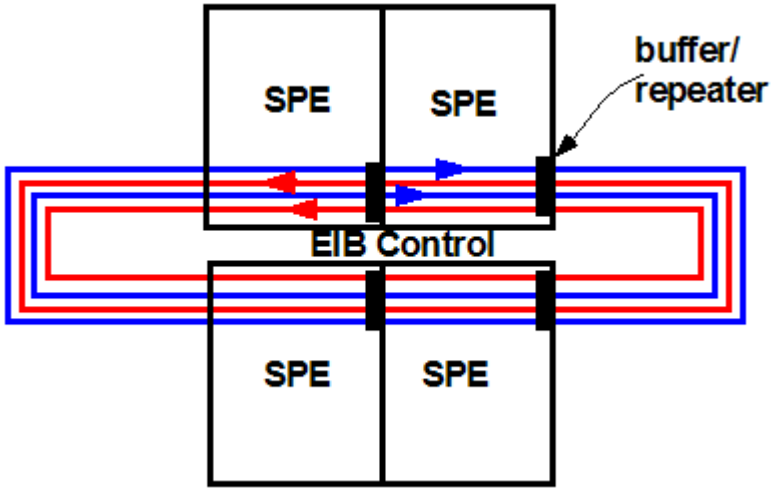

http://www.realworldtech.com/includes/images/articles/cell-9.gif

MIC

Memory Interface Controller

- Supports 2 Rambus XDR™ memory banks/channels
- Each bank 36 bits wide
- 2 XDR channels, 4 DRAM chips
  - ➢ 25.2 GB/s maximum bandwidth!
- MIC is asynchronous to processor and I/O interfaces
  - Speed-matching SRAM buffers/logic
  - Two clock domains
  - Transceiver training sequence req'd
  - Provides greater flexibility

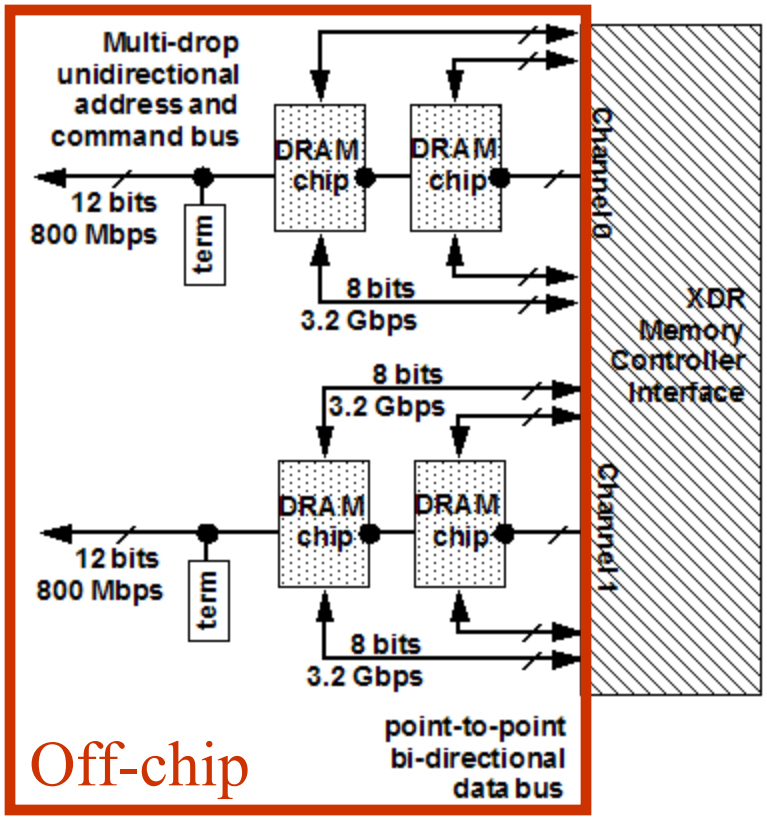

http://www.realworldtech.com/includes/images/articles/cell-10.gif

BIC

Bus Interface Controller

- Rambus FlexIO™ interface
- 44.8 GB/s outbound
- 32 GB/s inbound
- TX/RX asynchronous to proc./memory
- Supports multiple configurations w/o chip redesign or repackaging
- BIC provides asynch. interface b/n EIB and 2 I/Os
  - SPEs can snoop I/O activity via EIB
  - speed-matching SRAM buffers/logic
  - Three clock domains

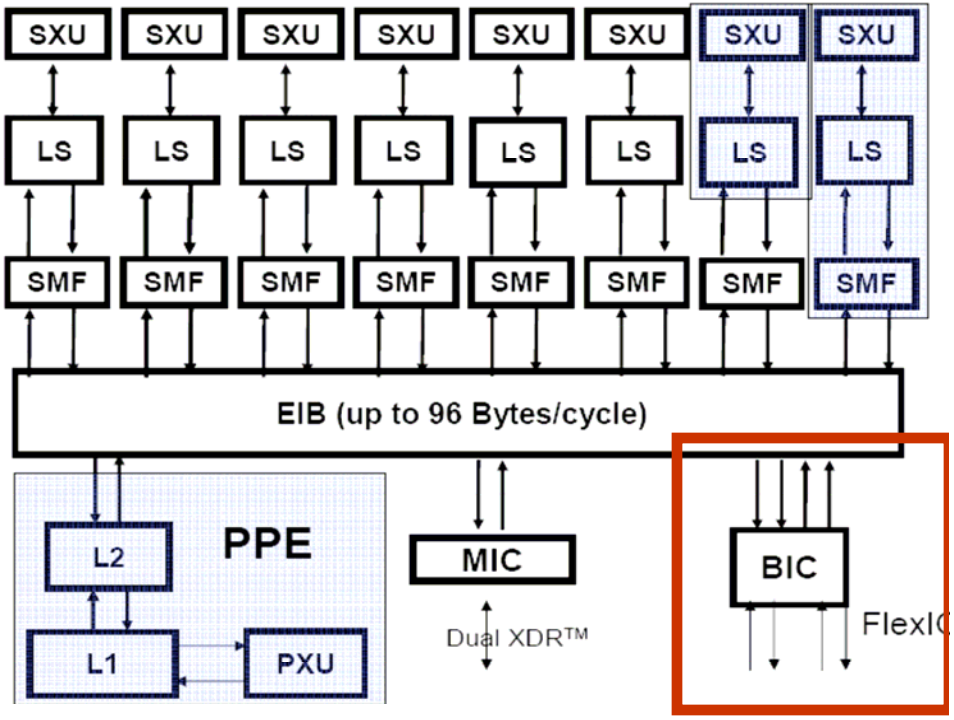

Pervasive Unit

- Contains all global logic for:
  - Basic functional operation
  - Lab debug
  - Manufacturing test
- Control for clock generation and distribution
- Power-on-reset (POR) engine
  - Contains debug mode allowing single-step, skipped, or out-of-order instruction sequences
- Performs performance analysis of all functional units on chip via trace/debug bus
- Provides logic used for programming and eFuses
  - Array repair and chip customization during manufacturing test

PMU

Power Management Unit

- Software controls to reduce chip power when full processing not needed
- OS has the power to:
  - Throttle...
  - Pause...
  - Stop...
  - ...single or multiple units or entire chip

http://www.nature.com/news/2004/040913/images/powerplant.jpg

TMU

Thermal Management Unit

- Two types of thermal sensors:
  - External linear diode
    - For controlling external cooling mechanisms
  - 10 digital thermal sensors (DTSs)
    - Distributed throughout chip
    - Temp of each element can be individually controlled by throttling
- Software controls TMU settings for each element
  - Four temperature settings for each element:
    - Throttling stops
    - Throttling starts
    - Element shut down
    - Chip clocks shut down

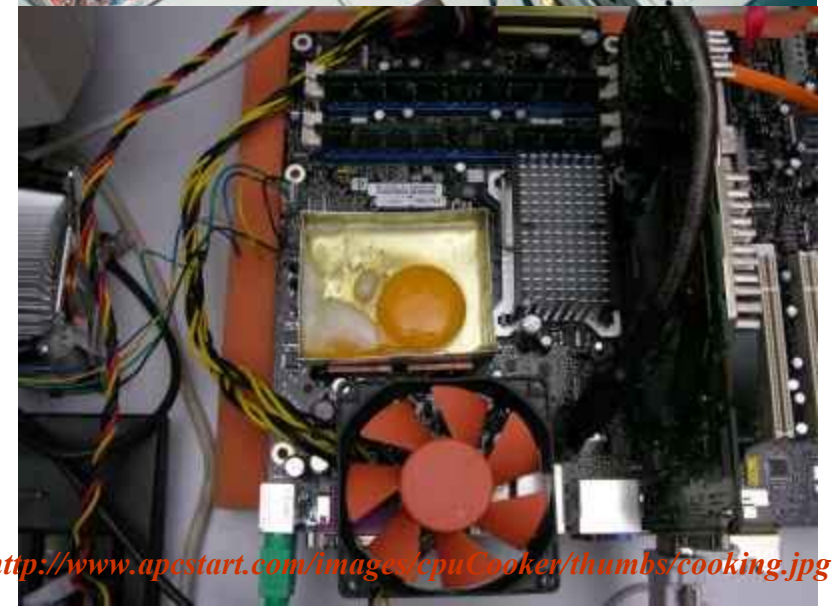

Digital Thermal Sensor

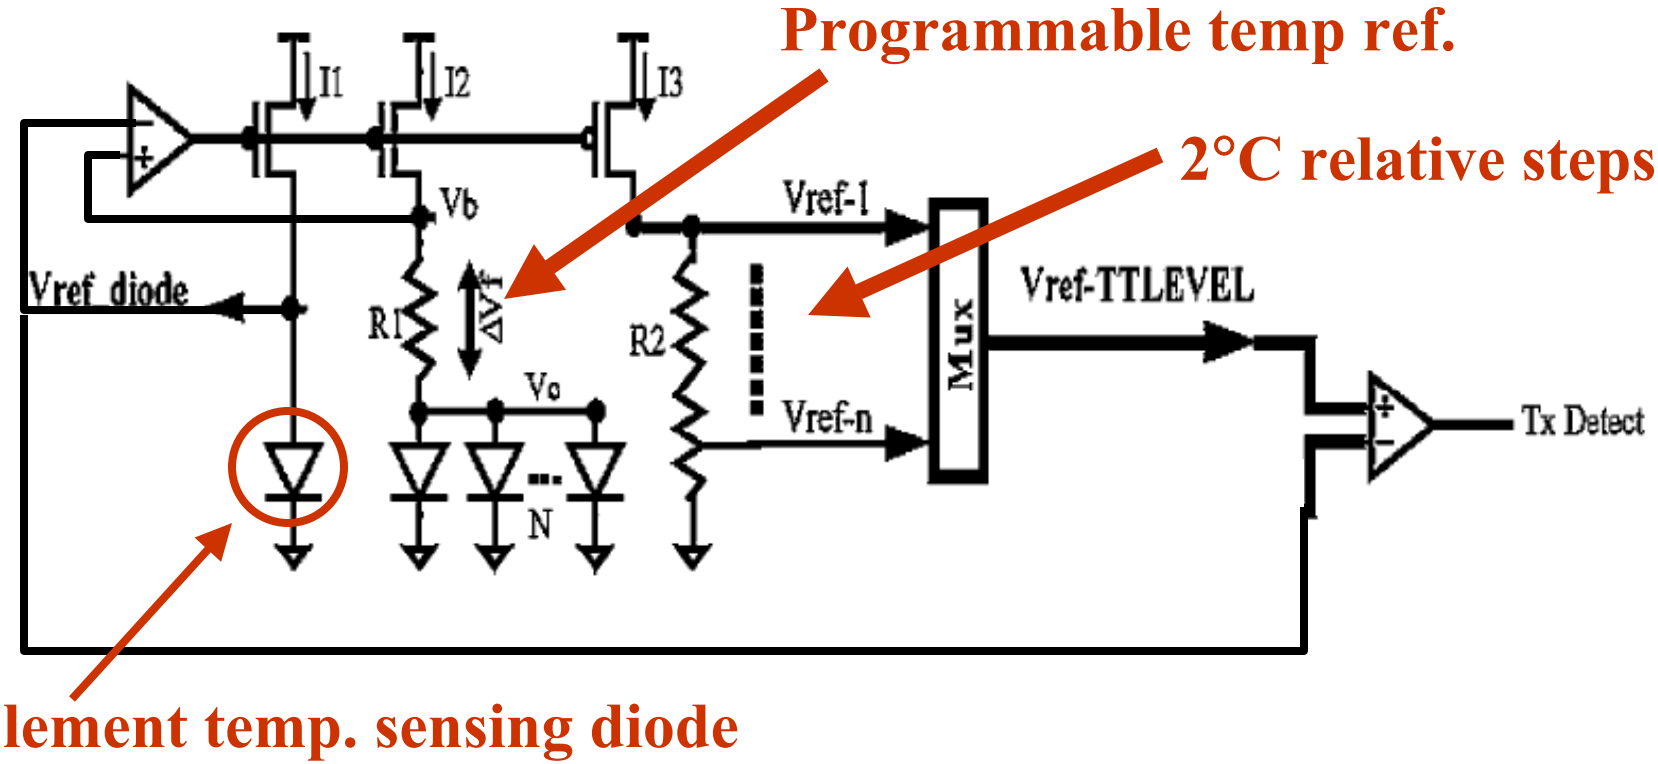

Technology Issues

- Additional HVT device
- Analog designs:
  - ↑ channel length
  - ↑ gate oxide thickness
  - Body-contacted devices
    - Vt control
  - Chip decoupling caps
- Interconnect
  - Too few layers ➔ ↑ area
  - Too many layers ➔ ↑ $$$

TABLE I
TECHNOLOGY PARAMETERS

FEOL	90 nm PD-SOI
Voltage	1.0 V
Thin oxide device	
Oxide thickness	10.5 A
Channel length	46 nm
Threshold voltages	Reg. V_T, High V_T, Super High V_T
Thick oxide device	
Oxide thickness	22.5 A
Channel length	140 nm
Threshold voltages	Reg. V_T
Decoupling cap oxide thickness	15A
BEOL	Local Interconnect plus 8LM
	5-1X layers, 1-2X layer, 2-6X layers

analog

Packing Issues

- Flip-chip PBGA
- Decoupling caps mounted on back against chip (close as possible to power pins)
  - Facilitated by fine-pitch mechanical drilling
  - Ensures robust mid-freq response
- Lid in thermal contact with die
- 3,349 C4s (solder bumps)

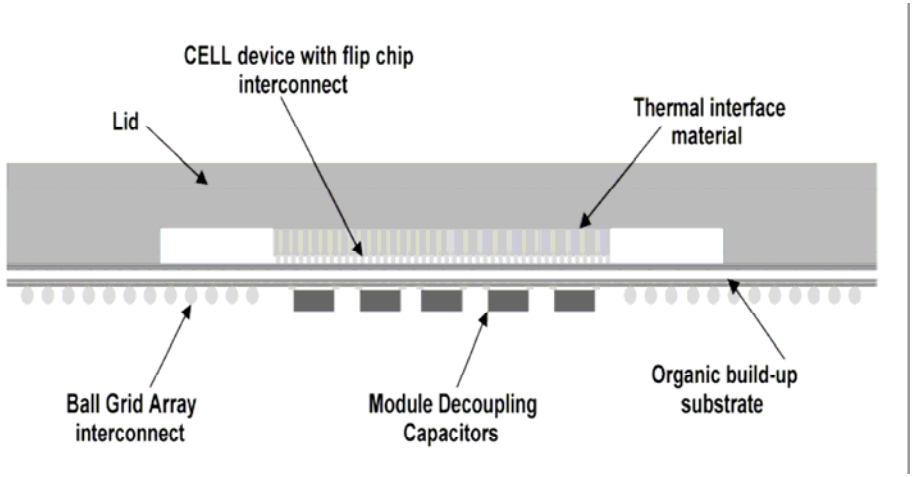

Clock Distribution

- Three distinct clock distribution systems
  - Processors
  - BIC
  - MIC
- Main high-freq grid
  - Covers 85% of chip
  - Buffered tunable trees
  - Final levels drive common grid
- Clock grids interleaved
  - Logic circuits can use all grids
- 850 individually tunable buffers (!)

- 3 clock grids
- 1,100 clock buffers
- 19.4 meters of metal!

Clock jitter and skew @ 16 GHz

- Avg.(?) Jitter
  - 12.7 ps peak-to-peak
  - 1.57 ps rms
- Min. Jitter
  - 9.48 ps peak-to-peak
  - 1.09 ps rms
- Skew
  - < 12 ps

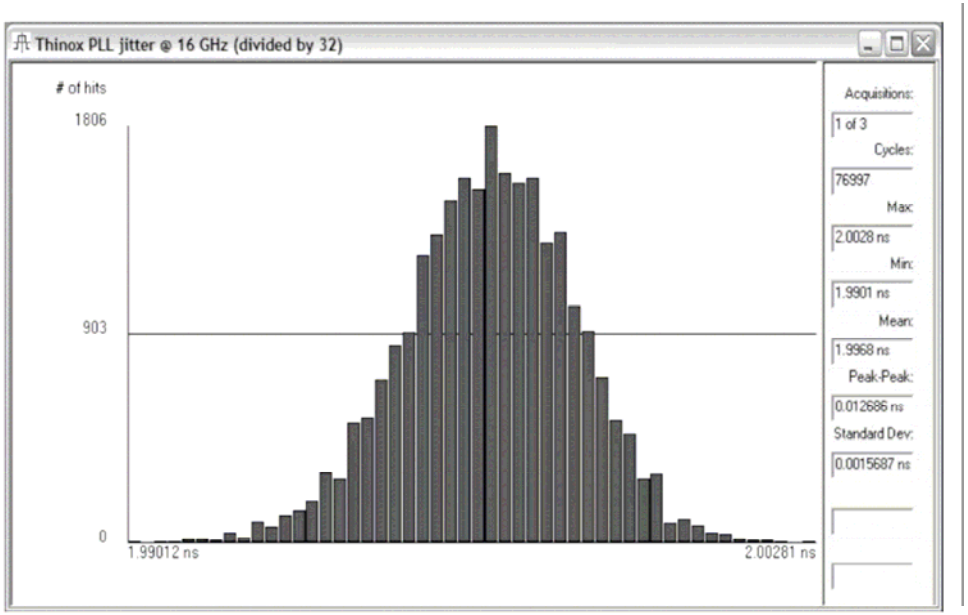

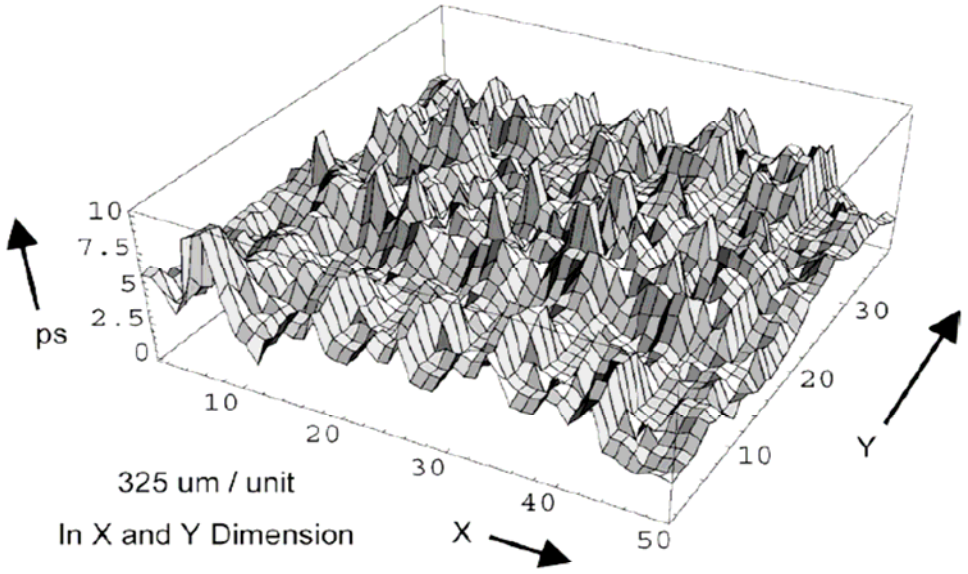

Fig. 13. Main clock distribution normalized arrival time.

Buffers/latches/FFs

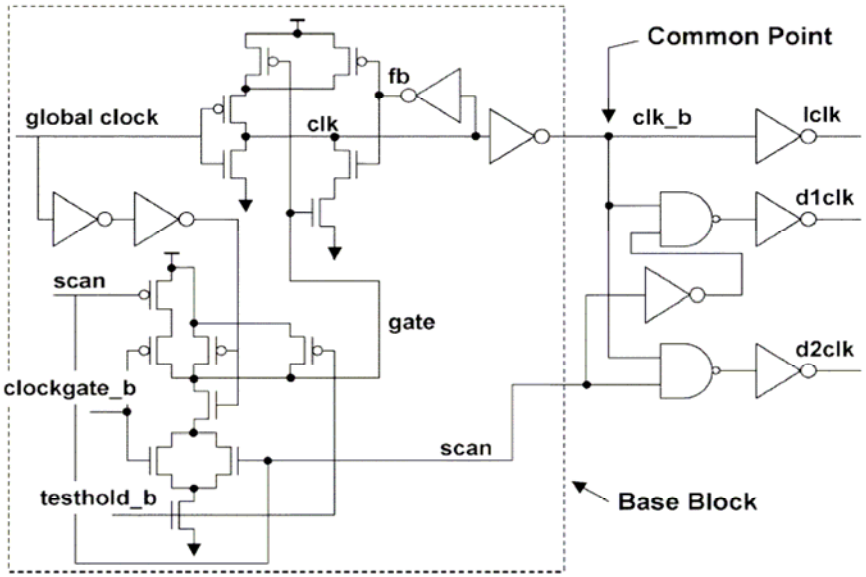

Fig. 14. Basic local clock components.

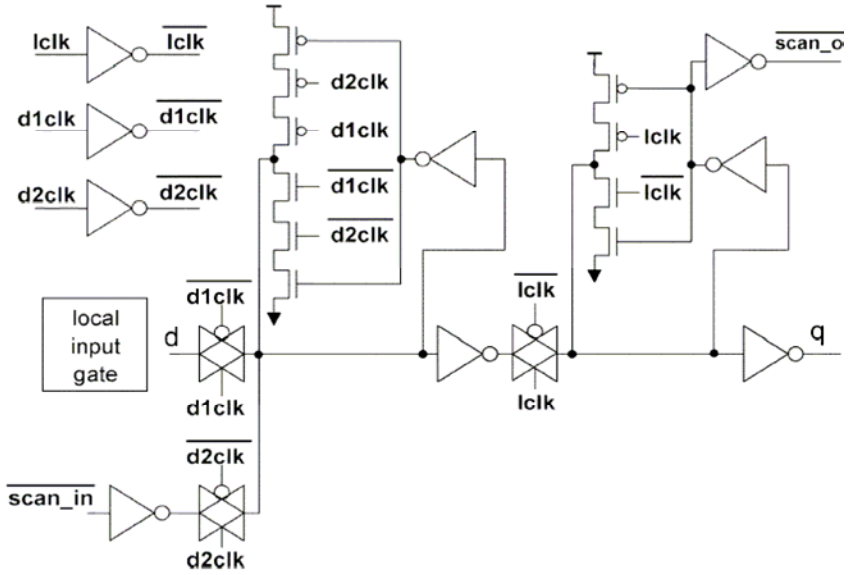

Fig. 15. Standard master–slave flip-flop.

- The point:
  - Minimize power consumption of clocked elements as a whole
  - Provide special high-speed constructs for critical paths
  - Single global clock, local clock splitters

Buffers/latches/FFs

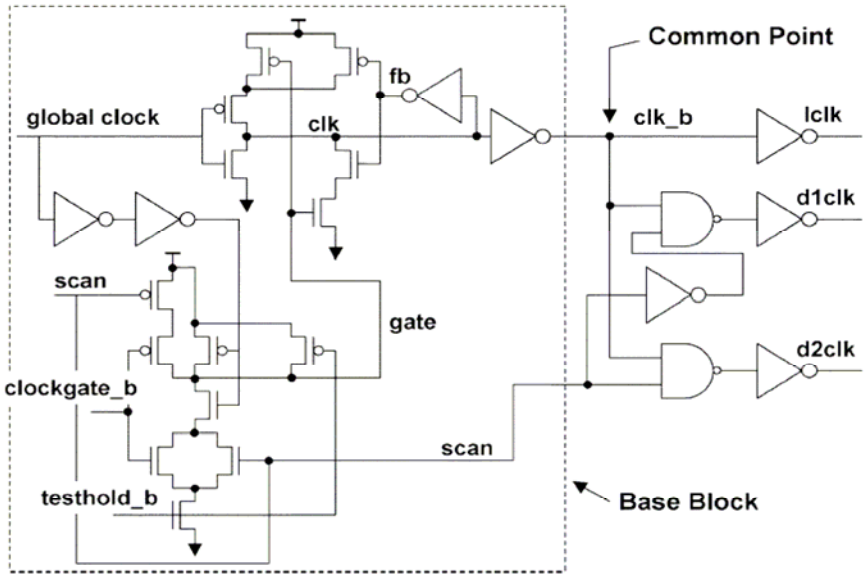

Fig. 14. Basic local clock components.

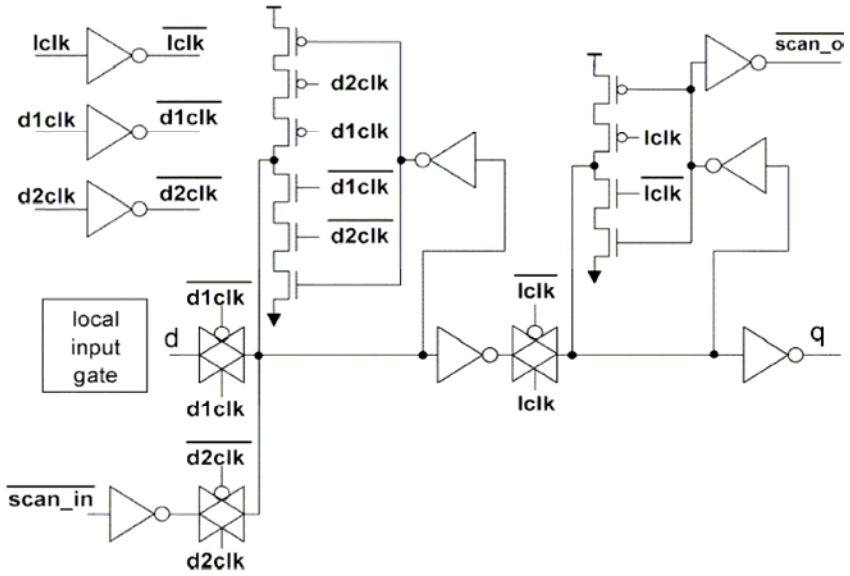

Fig. 15. Standard master–slave flip-flop.

- Support multiple operating modes
  - Local clock activate mode
  - Scan mode (scan)
  - Test mode (testhold_b)
  - State preservation

Buffers/latches/FFs

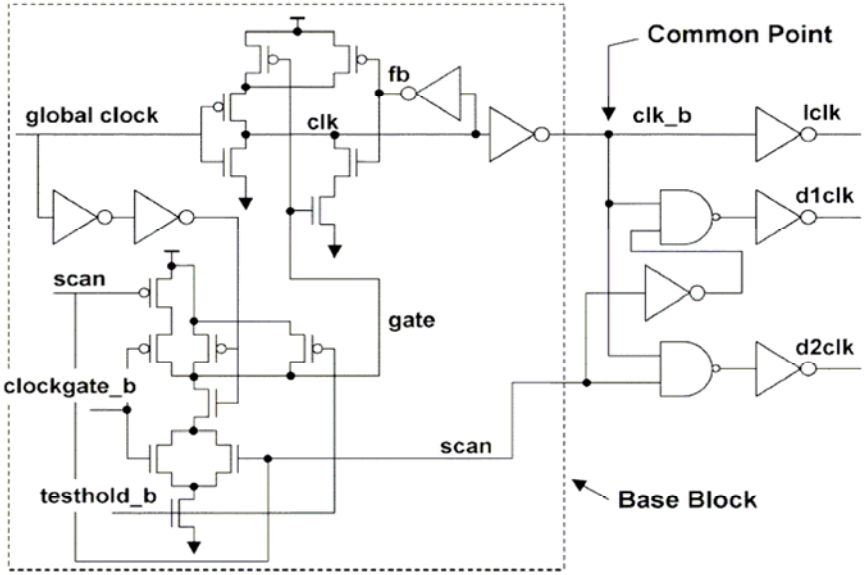

Fig. 14. Basic local clock components.

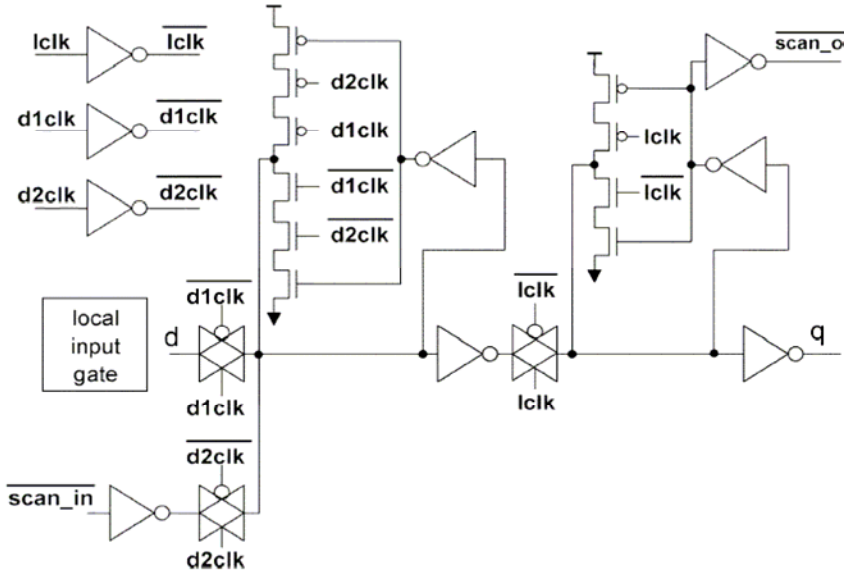

Fig. 15. Standard master–slave flip-flop.

- Support multiple operating modes
  - Local clock activate mode
  - Scan mode (scan)
  - Test mode (testhold_b)
  - Pulse mode (Fig. 17) ➔ "chicken-switch" used if race conditions occur

SRAM

- L1 cache:
  - 32-kB
  - 3-cycle latency
  - Parity checking
  - 64-byte write
  - 16-byte read
  - 16 cells / local BL
  - 8 LBLs / GBL
- L2 cache:
  - 512-kB
  - 1/8 active at a time
  - 280-bit reads supported
  - 2-cyc write / 3-cyc read
- SPE Local Stores
  - Four 64-kB arrays
  - 4-cyc write / 6-cyc read

Noise Analysis

- Macro-level
  - Transistor-level simulator
  - Netlist w/ parasitics extr. from macro layout
  - Static noise analysis
  - Macro noise abstract formed
    - Input noise tolerance
    - Input capacitance
    - Output resistance

- Unit/chip-level
  - Input:
    - Macro abstract
    - Timing analysis
  - Global netlist w/ parasitics extr. from global layout
  - Dynamic noise analysis using:
    - Resistors, caps, voltage sources
    - Timing information
  - Failure reported if noise level exceeds input tolerance

Power Analysis

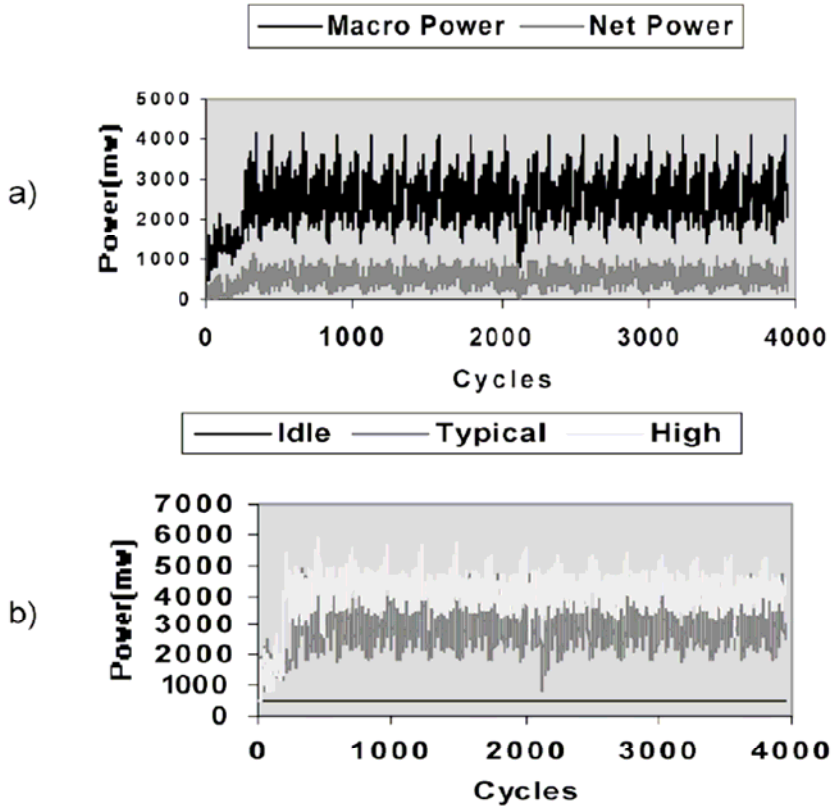

Fig. 25. (a) Macro and net power for a typical work load. (b) Total power.

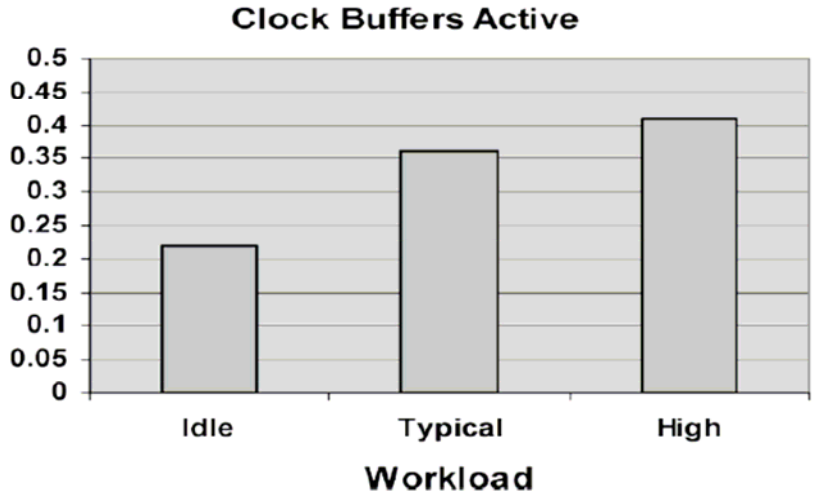

Fig. 26. Percent of local clock buffers active for idle, typical, and high-power workload.

Total Average Power: 60-80 Watts

Thermal Analysis

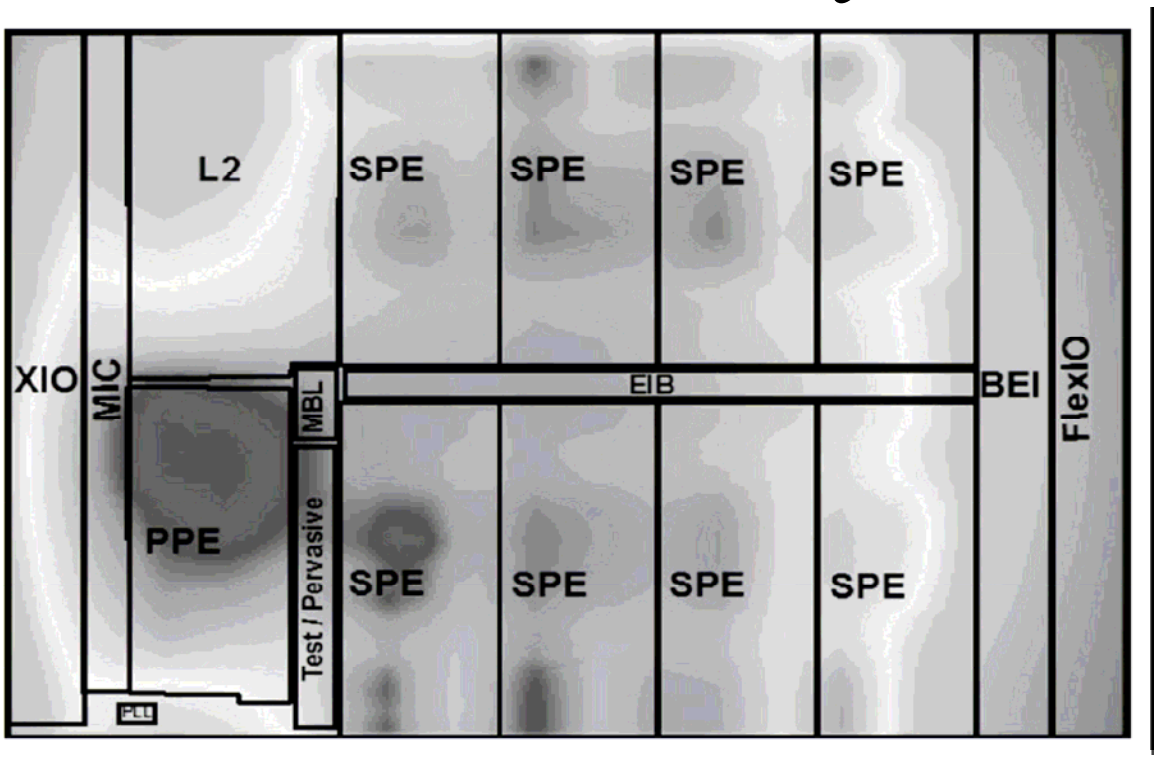

- Analytic modeling techniques based on
  - Classical steady-state thermal diffusion theory
  - Transient analysis under various workloads
- Considered ext. lateral thermal spreading through the package
- Analysis carried out early in design cycle
  - Affected floorplan design and thermal sensor design/placement

Conclusion

(*It works.)*

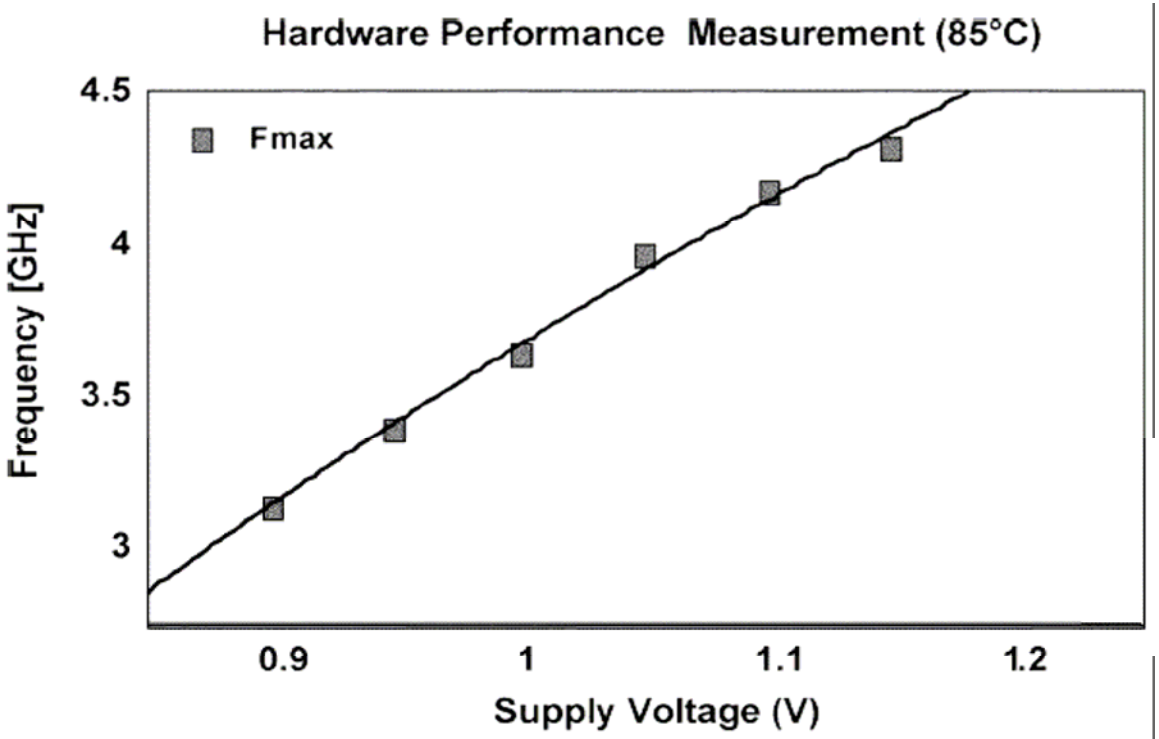

Fig. 28. First pass hardware results in the laboratory.

Except where stated, all figures and information on the CELL processor is from:
D. Pham, T. Aipperspach, D. Boerstler, M. Bolliger, R. Chaudhry, D. Cox, P. Harvey, P. M. Harvey, H. P. Hofstee, C. Johns, J. Kahle, A. Kameyama, J. Keaty, Y. Masubuchi, M. Pham, J. Pille, S. Posluszny, M. Riley, D. Stasiak, M. Suzuoki, O. Takahashi, J. Warnock, S. Weitzel, D. Wendel, and K. Yazawa, "Overview of the Architecture, Circuit Design, and Physical Implementation of a First-Generation Cell Processor," *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 1, pages 179-196, January 2006.

http://en.wikipedia.org/wiki/Image:ME0000570927_2.jpg

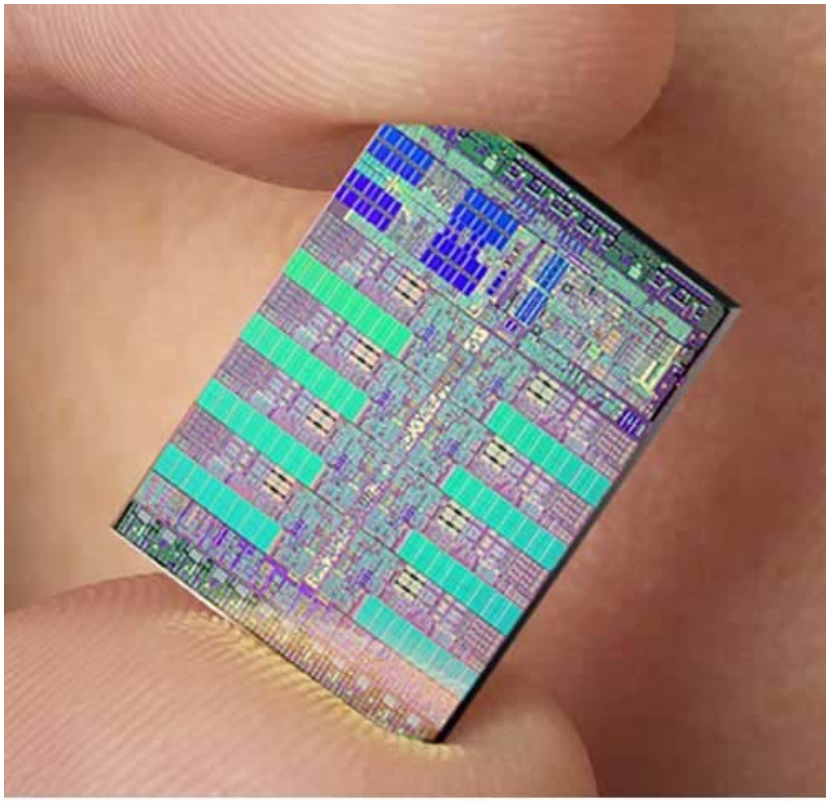

Implemented on a 90-nm silicon-on-insulator process with eight levels of copper interconnect, the CELL processor can deliver a single-precision compute throughput of over 256 GFLOPS when running at 4 GHz.

http://www.elecdesign.com/Files/29/9748/Figure_02.jpg